

# Towards a Benchmark and a Comparison Framework for Combinatorial Interaction Testing of Software Product Lines

Roberto E.  
Lopez-Herrejon  
Johannes Kepler University  
Linz, Austria  
rlopez@jku.at

Javier Ferrer  
University of Málaga, Spain  
ferrer@lcc.uma.es

Francisco Chicano  
University of Málaga, Spain  
chicano@lcc.uma.es

Evelyn Nicole Haslinger  
Johannes Kepler University  
Linz, Austria  
evelyn.haslinger@jku.at

Alexander Egyed  
Johannes Kepler University  
Linz, Austria  
alexander.egyed@jku.at

Enrique Alba  
University of Málaga, Spain  
eat@lcc.uma.es

## ABSTRACT

As Software Product Lines (SPLs) are becoming a more pervasive development practice, their effective testing is becoming a more important concern. In the past few years many SPL testing approaches have been proposed, among them, are those that support Combinatorial Interaction Testing (CIT) whose premise is to select a group of products where faults, due to feature interactions, are more likely to occur. Many CIT techniques for SPL testing have been put forward; however, no systematic and comprehensive comparison among them has been performed. To achieve such goal two items are important: a common benchmark of feature models, and an adequate comparison framework. In this research-in-progress paper, we propose 19 feature models as the base of a benchmark, which we apply to three different techniques in order to analyze the comparison framework proposed by Perrouin et al. We identify the shortcomings of this framework and elaborate alternatives for further study.

## Keywords

Combinatorial Interaction Testing, Software Product Lines, Pairwise Testing, Feature Models

## 1. INTRODUCTION

A *Software Product Line (SPL)* is a family of related software systems, which provide different feature combinations [31]. The effective management and realization of *variability* – the capacity of software artifacts to vary [35] – can lead to substantial benefits such as increased software reuse, faster product customization, and reduced time to market.

Systems are being built, more and more frequently, as SPLs rather than individual products because of several technological and marketing trends. This fact has created an increasing need for testing approaches that are capable of coping with large numbers of feature combinations that characterize SPLs. Many testing alternatives have been put forward [9, 12, 14, 22]. Salient among them are those that support *Combinatorial Interaction Testing (CIT)* whose premise is to select a group of products where faults, due to feature interactions, are more likely to occur. In particular, most of the work has focused on pairwise testing whereby the interactions of two features are considered [7, 15, 17, 19, 23, 28, 30]. With all these pairwise testing approaches available the question now is: how do they compare? To answer this question, two items are necessary: a common benchmark of feature models, and an adequate comparison framework. In this research-in-progress paper, we propose a set of 19 feature models as a base of the comparison benchmark. In addition, we use these feature models to illustrate Perrouin et al.'s comparison framework [29]. We identify some shortcomings, and elaborate alternatives for further study.

The organization of the paper is as follows. In Section 2 we present the basic background on feature models. Section 3 describes the basic terminology of CIT and how it is applied to SPLs. Section 4 presents the list of feature models that we proposed as basic benchmark. Section 5 summarizes and illustrates Perrouin et al.'s comparison framework. Section 6 sketches the three CIT algorithms used to illustrate both the benchmark and the comparison framework. Section 7 presents the results of our evaluation and its analysis. Section 8 briefly summarizes the related work and Section 9 outlines the conclusions and future work.

## 2. FEATURE MODELS AND RUNNING EXAMPLE

Feature models have become a *de facto* standard for modelling the common and variable features of an SPL and their relationships collectively forming a tree-like structure. The nodes of the tree are the features, which are depicted as labelled boxes, and the edges represent the relationships




Figure 1: Graph Product Line Feature Model

among them. Thus, a feature model denotes the set of feature combinations that the products of an SPL can have [21].

Figure 1 shows the feature model of our running example, the *Graph Product Line (GPL)*, a standard SPL of basic graph algorithms that has been extensively used as a case study in the product line community [24]. A product has feature **GPL** (the root of the feature model) which contains its core functionality, and a driver program (**Driver**) that sets up the graph examples (**Benchmark**) to which a combination of graph algorithms (**Algorithms**) are applied. The types of graphs (**GraphType**) can be either directed (**Directed**) or undirected (**Undirected**), and can optionally have weights (**Weight**). Two graph traversal algorithms (**Search**) are available: either Depth First Search (**DFS**) or Breadth First Search (**BFS**). A product must provide at least one of the following algorithms: numbering of nodes in the traversal order (**Num**), connected components (**CC**), strongly connected components (**SCC**), cycle checking (**Cycle**), shortest path (**Shortest**), minimum spanning trees with Prim’s algorithm (**Prim**) or Kruskal’s algorithm (**Kruskal**).

In a feature model, each feature (except the root) has one parent feature and can have a set of child features. Notice here that a child feature can only be included in a feature combination of a valid product if its parent is included as well. The root feature is always included. There are four kinds of feature relationships:

- *Mandatory features* are depicted with a filled circle. A mandatory feature is selected whenever its respective parent feature is selected. For example, features **Driver** and **GraphType**.
- *Optional features* are depicted with an empty circle. An optional feature may or may not be selected if its respective parent feature is selected. An example is feature **Weight**.
- *Exclusive-or relations* are depicted as empty arcs crossing over a set of lines connecting a parent feature with its child features. They indicate that exactly one of the features in the exclusive-or group must be selected whenever the parent feature is selected. For example, if feature **Search** is selected, then either feature **DFS** or feature **BFS** must be selected.
- *Inclusive-or relations* are depicted as filled arcs crossing over a set of lines connecting a parent feature with its child features. They indicate that at least one of

the features in the inclusive-or group must be selected if the parent is selected. If for instance, feature **Algorithms** is selected then at least one of the features **Num**, **CC**, **SCC**, **Cycle**, **Shortest**, **Prim**, or **Kruskal** must be selected.

Besides the parent-child relations, features can also relate across different branches of the feature model with *Cross-Tree Constraints (CTCs)*. Figure 1 shows the CTCs of GPL. For instance, **Cycle** **requires** **DFS** means that whenever feature **Cycle** is selected, feature **DFS** must also be selected. As another example, **Prim** **excludes** **Kruskal** means that both features cannot be selected at the same time in any product. These constraints as well as those implied by the hierarchical relations between features are usually expressed and checked using propositional logic, for further details refer to [6]. Now we present the basic definitions on which SPL testing terminology is defined in the next section.

**Definition 1.** *Feature List (FL) is the list of features in a feature model.*

The FL for the GPL feature model is [GPL, Driver, Benchmark, GraphType, Directed, Undirected, Weight, Search, DFS, BFS, Algorithms, Num, CC, SCC, Cycle, Shortest, Prim, Kruskal].

**Definition 2.** *A feature set, also called product in an SPL, is a 2-tuple  $[sel, \overline{sel}]$  where  $sel$  and  $\overline{sel}$  are respectively the set of selected and not-selected features of a member product<sup>1</sup>. Let  $FL$  be a feature list, thus  $sel, \overline{sel} \subseteq FL$ ,  $sel \cap \overline{sel} = \emptyset$ , and  $sel \cup \overline{sel} = FL$ . The terms  $p.sel$  and  $p.\overline{sel}$  respectively refer to the set of selected and not-selected features of product  $p$ .*

**Definition 3.** *A feature set  $fs$  is valid in feature model  $fm$ , iff  $fs$  does not contradict any of the constraints introduced by  $fm$ . We will denote with  $FS$  the set of valid feature sets for a feature model (we omit the feature model in the notation).*

For example, the feature set  $fs0 = \{[GPL, Driver, GraphType, Weight, Algorithms, Benchmark, Undirected, Prim], [Search, Directed, DFS, BFS, Num, CC, SCC, Cycle, Shortest, Kruskal]\}$  is valid. As another example, a feature set with features **DFS** and **BFS** would not be valid because it violates the constraint of the exclusive-or relation which establishes that these two features cannot appear selected together in the same feature set. The GPL feature model denotes 73 valid feature sets, some of them depicted in Table 1, where selected features are ticked (✓) and unselected features are empty.

**Definition 4.** *A feature  $f$  is a core feature if it is selected in all the valid feature sets of a feature model  $fm$ , and is a variant feature if it is selected in some of the feature sets.*

For example **GPL**, **Driver**, **Benchmark**, **GraphType** and **Algorithms** are core features and the remaining ones are variant features.

### 3. COMBINATORIAL INTERACTION TESTING FOR SOFTWARE PRODUCT LINES

Combinatorial Interaction Testing (CIT) is a testing approach that constructs samples to drive the systematic testing of software system configurations [8]. When applied to

<sup>1</sup>Definition based on [6].

Table 1: Sample Feature Sets of GPL

FS	GPL	Dri	Gtp	W	Se	Alg	B	D	U	DFS	BFS	N	CC	SCC	Cyc	Sh	Prim	Kru
fs0	✓	✓	✓	✓	✓	✓	✓		✓								✓	
fs1	✓	✓	✓	✓	✓	✓	✓		✓	✓			✓					✓
fs2	✓	✓	✓	✓	✓	✓	✓	✓		✓					✓			
fs3	✓	✓	✓	✓	✓	✓	✓	✓			✓	✓				✓		
fs4	✓	✓	✓	✓	✓	✓	✓	✓		✓			✓	✓	✓	✓		
fs5	✓	✓	✓	✓	✓	✓	✓		✓	✓		✓	✓		✓		✓	
fs6	✓	✓	✓	✓	✓	✓	✓		✓		✓	✓	✓					✓
fs7	✓	✓	✓	✓	✓	✓	✓		✓	✓		✓	✓		✓			

Driver (Dri), GraphType (Gtp), Weight (W), Search (Se), Algorithms (Alg), Benchmark (B), Directed (D), Undirected (U), Num (N), Cycle (Cyc), Shortest (Sh), Kruskal (Kr).

SPL testing, the idea is to select a representative subset of products where interaction errors are more likely to occur rather than testing the complete product family [8]. In this section we provide the basic terminology of CIT within the context of SPLs.

**Definition 5.** A  $t$ -set  $ts$  is a 2-tuple  $[sel, \overline{sel}]$  representing a partially configured product, defining the selection of  $t$  features of the feature list  $FL$ , i.e.  $ts.sel \cup ts.\overline{sel} \subseteq FL \wedge ts.sel \cap ts.\overline{sel} = \emptyset \wedge |ts.sel \cup ts.\overline{sel}| = t$ . We say  $t$ -set  $ts$  is covered by feature set  $fs$  iff  $ts.sel \subseteq fs.sel \wedge ts.\overline{sel} \subseteq fs.\overline{sel}$ .

**Definition 6.** A  $t$ -set  $ts$  is valid in a feature model  $fm$  if there exists a valid feature set  $fs$  that covers  $ts$ . The set of all valid  $t$ -sets for a feature model is denoted with  $TS^2$ .

**Definition 7.** A  $t$ -wise covering array  $tCA$  for a feature model  $fm$  is a set of valid feature sets that covers all valid  $t$ -sets denoted by  $fm^3$ . We also use the term test suite to refer to a covering array.

Let us illustrate these concepts for pairwise testing, meaning  $t=2$ . From the feature model in Figure 1, a valid 2-set is  $[\{\text{Driver}\}, \{\text{Prim}\}]$ . It is valid because the selection of feature **Driver** and the non-selection of feature **Prim** do not violate any constraints. As another example, the 2-set  $[\{\text{Kruskal}, \text{DFS}\}, \emptyset]$  is valid because there is at least one feature set, for instance **fs1** in Table 1, where both features are selected. The 2-set  $[\emptyset, \{\text{SCC}, \text{CC}\}]$  is also valid because there are valid feature sets that do not have any of these features selected, for instance feature sets **fs0**, **fs1**, and **fs3**. Notice however that the 2-set  $[\emptyset, \{\text{Directed}, \text{Undirected}\}]$  is not valid. This is because feature **GraphType** is present in all the feature sets (mandatory child of the root) so either **Directed** or **Undirected** must be selected. In total, our running example has 418 valid 2-sets.

Based on Table 1, the three valid 2-sets just mentioned above are covered as follows. The 2-set  $[\{\text{Driver}\}, \{\text{Prim}\}]$  is covered by feature sets **fs1**, **fs2**, **fs3**, **fs4**, **fs6**, and **fs7**. Similarly, the 2-set  $[\{\text{Kruskal}, \text{DFS}\}, \emptyset]$  is covered by feature set **fs1**, and  $[\emptyset, \{\text{SCC}, \text{CC}\}]$  is covered by feature sets **fs0**, **fs2**, and **fs3**.

## 4. BASIC BENCHMARK

We propose the use of 19 *realistic* feature models as a basis for a comparison benchmark. By realistic we mean that these models meet three basic requirements:

<sup>2</sup>We also omit here the feature model in the notation for the sake of clarity.

<sup>3</sup>Definition based on [20].

1. *Available Source Code.* Because the ultimate goal of this line of research is to evaluate the effectiveness of the testing approaches, it is thus of the utmost importance that the source code associated to the proposed feature models be available in a complete form, although perhaps not be thoroughly documented.
2. *Explicit Feature Model.* We consider feature models that are explicitly provided by the SPL authors. This requirement is to prevent any misunderstandings or omissions that any techniques to reverse engineering feature models from other artifacts can potentially have.
3. *Plausible number of products.* It does not take many features to create feature models with a huge number of potential products. We arbitrarily chose two million as the maximum number of products denoted by the feature models in the benchmark. We would argue this is a reasonable number of products that a large company or open source community could potentially maintain and most importantly thoroughly test.

In order to find the feature models that meet these criteria we searched proceedings from SPL-related venues such as SPLC, VAMOS, ICSE, ASE, and FSE published over the last five years. In addition, we consulted the following websites and repositories: SPL Conqueror [34], FeatureHouse [2], SPL2go [4], and SPLOT [3]<sup>4</sup>. Table 2 summarizes the feature models used in our evaluation. It shows the number of features, number of products, and their application domain with the reference from where they were obtained.

We should point out that some of the pairwise testing approaches identified and mentioned in Section 8 already use some examples from the SPLOT website; however, to the best of our knowledge, the criteria for their feature models selection is not specified precisely. In our experience with this repository, based on the information provided by the model authors on the SPLOT website itself, we either were not able to trace the code sources of the feature models or we found semantic mistakes in them.

We should stress that this list of feature models is by no means complete. Our expectation, as a result of this paper, is that the SPL community proposes new feature models to add or remove to this benchmark, perhaps filling in details that were not found by our search, and adding or refining our selection criteria.

<sup>4</sup>Search performed during August-September 2013.

Feature Model	NF	NP	Domain
Apache	10	256	web server [34]
argo-uml-spl	11	192	UML tool [1]
BDB*	117	32	database [2]
BDBFootprint	9	256	database [34]
BDBMemory	19	3,840	database [34]
BDBPerformance	27	1,440	database [34]
Curl	14	1024	data transfer [34]
DesktopSearcher	22	462	file search [4]
fame_dbms_fm	20	320	database [4]
gpl	18	73	graph algorithms [24]
LinkedList	27	1,344	data structures [34]
LLVM	12	1,024	compiler library [34]
PKJab	12	72	messenger [34]
Prevayler	6	32	object persistence [34]
SensorNetwork	27	16,704	networking [34]
TankWar	37	1,741,824	game [2]
Wget	17	8,192	file retrieval [34]
x264	17	2,048	video encoding [34]
ZipMe	8	64	data compression [34]

NF: Number of Features, NP: Number of Products,  
\*BDB prefix standards for Berkeley database.

Table 2: Feature Models Summary

## 5. COMPARISON FRAMEWORK

In this section we present the four metrics that constitute the framework proposed by Perrouin *et al.* for the comparison of pairwise testing approaches for SPLs [29]. We define them based on the terminology presented in Sections 2 and 3. For the following metric definitions, let  $tCA$  be a  $t$ -wise covering array of feature model  $fm$ . The corresponding equations are shown in Figure 2.

**Metric 1. Test Suite Size** is the number of feature sets selected in a covering array for a feature model, shown in Equation (1).

**Metric 2. Performance** is the time required for an algorithm to compute a covering array.

**Metric 3. Test Suite Similarity.** This metric is defined based on Jaccard’s similarity index and applied to variant features. Let  $FM$  be the set of all possible feature models,  $fs$  and  $gs$  be two feature sets in  $FS$ , and  $var : FS \times FM \rightarrow FL$  be an auxiliary function that returns the selected variant features of a feature set according to a FM. The similarity index of two feature sets is thus defined in Equation (2), and the similarity value for the entire covering array is defined by Equation (3).

It should be noted here that the second case of the similarity index, when there are no variant features on both feature sets, is not part of the original proposed comparison framework [29]. We added this term because in our search we found feature sets formed only with core features.

**Metric 4. Tuple Frequency.** Let  $occurrence : TS \times 2^{FS} \rightarrow \mathbb{N}$  be an auxiliary function that counts the occurrence of a  $t$ -set (a tuple of  $t$  elements) in all the feature sets of a covering array of a single feature model. The metric is defined in Equation (4).

The first two metrics are the standard measurements used for comparison between different testing algorithms, not only within the SPL domain. To the best of our understanding, the intuition behind the Test Suite Similarity is that the more dissimilar (value close to 0) the feature sets are, the higher chances to detect any faulty behaviour when the corresponding  $t$ -wise tests are instrumented and performed.

Along the same lines, the rationale behind tuple frequency is that by reducing this number, the higher the chances of reducing the repetition of executions of  $t$ -wise tests.

Let us provide some examples for the latter two metrics for the case of pairwise testing,  $t=2$ . Consider for instance, feature sets  $fs0$ ,  $fs1$ ,  $fs2$  and  $fs7$  from Table 1. The variant features in those feature sets are:

$$var(fs0, gpl) = \{Undirected, Weight, Prim\}$$

$$var(fs1, gpl) = \{Undirected, Weight, Search, DFS, Connected, Kruskal\}$$

$$var(fs2, gpl) = \{Directed, Search, DFS, Number, Cycle\}$$

$$var(fs7, gpl) = \{Undirected, Weight, Search, DFS, Connected, Number, Cycle\}$$

An example is the similarity value between feature sets  $fs0$  and  $fs2$ , that is  $Sim(fs0, fs2, gpl) = 0/8 = 0.0$ . The value is zero because those two feature sets do not have any selected variant features in common. Now consider  $Sim(fs1, fs7, gpl) = 5/8 = 0.625$  which yields a high value because those feature sets have the majority of their selected features in common.

For sake of illustrating the Tuple Frequency metric, let us assume that the set of feature sets in Table 1 is a 2-wise covering array of GPL denoted as  $tCA_{gpl}^5$ . For example, the 2-set  $ts0 = [\{Driver\}, \{Prim\}]$  is covered by feature sets  $fs1$ ,  $fs2$ ,  $fs3$ ,  $fs4$ ,  $fs6$ , and  $fs7$ . Thus, its frequency is equal to  $occurrence(ts0, tCA_{gpl})/8 = 6/8 = 0.75$ . As another example, the 2-set  $ts1 = [\{Kruskal, DFS\}, \emptyset]$  is covered by feature set  $fs1$ . Thus its frequency is equal to  $occurrence(ts1, tCA_{gpl})/8 = 1/8 = 0.125$ .

Next we present the three algorithms that we used to assess this comparison framework on the feature models of the proposed benchmark.

## 6. ALGORITHMS OVERVIEW

In this section we briefly describe the three testing algorithms we used in our study.

### 6.1 CASA Algorithm

CASA is a simulated annealing algorithm that was designed to generate  $n$ -wise covering arrays for SPLs [17]. CASA relies on three nested search strategies. The outermost search performs one-sided narrowing, pruning the potential size of the test suite to be generated by only decreasing the upper bound. The mid-level search performs a binary search for the test suite size. The innermost search strategy is the actual simulated annealing procedure, which tries to find a pairwise test suite of size  $N$  for feature model  $FM$ . For more details on CASA please refer to [17].

### 6.2 Prioritized Genetic Solver

The *Prioritized Genetic Solver (PGS)* is an evolutionary approach proposed by Ferrer *et al.* [16] that constructs a test suite taking into account priorities during the generation.

<sup>5</sup>There are 24 2-wise pairs, out of the 418 pairs that GPL contains, which are not covered.

Figure 2: Framework Metrics Summary

$$TestSuiteSize = |tCA| \quad (1)$$

$$Similarity(fs, gs, fm) = \begin{cases} \frac{|var(fs, fm) \cap var(gs, fm)|}{|var(fs, fm) \cup var(gs, fm)|} & \text{if } var(fs, fm) \cup var(gs, fm) \neq \emptyset \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

$$TestSuiteSimilarity(tCA, fm) = \frac{\sum_{fsi \in tCA} \sum_{fsj \in tCA} Similarity(fsi, fsj, fm)}{|tCA|^2} \quad (3)$$

$$TupleFrequency(ts, tCA) = \frac{occurrence(ts, tCA)}{|tCA|} \quad (4)$$

PGS is a constructive genetic algorithm that adds one new product to the partial solution in each iteration until all pairwise combinations are covered. In each iteration the algorithm tries to find the product that adds the most coverage to the partial solution. This paper extends and adapts PGS for SPL testing. PGS has been implemented using jMetal [13], a Java framework aimed at the development, experimentation, and study of metaheuristics for solving optimization problems. For further details on PGS, please refer to [16].

### 6.3 ICPL

ICPL is a greedy approach to generate  $n$ -wise test suites for SPLs, which has been introduced by Johansen *et al.* [20]<sup>6</sup>. It is basically an adaptation of Chvátal’s algorithm to solve the set cover problem. First, the set  $TS$  of all valid  $t$ -sets that need to be covered is generated. Next, the first feature set (product)  $fs$  is computed by greedily selecting a subset of  $t$ -sets in  $TS$  that constitute a valid product in the input feature model and added to the (initially empty) test suite  $tCA$ . Henceforth, all  $t$ -sets that are covered by product  $fs$  are removed from  $TS$ . ICPL then proceeds to generate products and adds them to the test suite  $tCA$  until  $TS$  is empty, i.e. all valid  $t$ -sets are covered by at least one product. To increase ICPLs performance Johansen *et al.* made several enhancements to the algorithm, for instance they parallelized the data independent processing steps. For further details on ICPL please refer to [20].

## 7. EVALUATION

In this section we present the evaluation of the benchmark using the comparison framework for the case of pairwise testing. We described the statistical analysis performed and the issues found. All the data and code used in our analysis is available on [25].

### 7.1 Experimental Set Up

The three algorithms, CASA, PGS and ICPL, are non-deterministic. For this reason we performed 30 independent runs for a meaningful statistical analysis. All the executions were run in a cluster of 16 machines with Intel Core2 Quad

processors Q9400 (4 cores per processor) at 2.66 GHz and 4 GB memory running Ubuntu 12.04.1 LTS and managed by the HT Condor 7.8.4 cluster manager. Since we have 3 algorithms and 19 feature models the total number of independent runs is  $3 \cdot 19 \cdot 30 = 1,710$ . Once we obtained the resulting test suites we applied the metrics defined in Section 5 and we report summary statistics of these metrics. In order to check if the differences between the algorithms are statistically significant or just a matter of chance, we applied the Wilcoxon rank-sum [33] test. In order to properly interpret the results of statistical tests, it is always advisable to report effect size measures. For that purpose, we have also used the non-parametric effect size measure  $\hat{A}_{12}$  statistic proposed by Vargha and Delaney [37], as recommended by Arcuri and Briand [5].

### 7.2 Analysis

The first thing we noticed is the inadequacy of the Tuple Frequency metric for our comparison purposes. Its definition, as shown in Equation (4), applies on a per tuple (i.e.  $t$ -set) basis. In other words, given a test suite, we should compute this metric for each tuple and we should provide the histogram of the tuple frequency. This is what the authors of [29] do. This means we should show 1,710 histograms, one per test suite which evidently is not a viable option to aggregate the information of this metric. An alternative is also presented in [29]. It consists on using the average of the tuple frequencies in a test suite taking into account all the tuples. Unfortunately we found that this average says nothing about the test suite. By using counting arguments we show that this average depends only on the number of features and the number of valid tuples of the model (i.e. the average is the same for all the test suites associated to a feature model), hence it is not suited to assess the quality of the test suite. The proof is presented in the Appendix. In the following we omit any information related to the Tuple Frequency metric and defer to our future work to study how to aggregate it.

In order to assess whether there was a correlation between the feature metrics, we calculated the Spearman rank’s correlation coefficient for each pair of metrics. Table 3 shows the results obtained plus the correlation values with the number of products and the number of features of the FMs

<sup>6</sup>ICPL stands for “*ICPL Covering array generation algorithm for Product Lines*”.

Table 3: Spearman’s correlation coefficients of all models and algorithms.

	Products	Features	TSSize	Performance	Similarity
Products	1	0.495	0.717	0.169	-0.015
Features	0.495	1	0.537	0.336	0.180
TSSize	0.717	0.537	1	0.280	-0.106
Performance	0.169	0.336	0.280	1	0.440
Similarity	-0.015	0.180	-0.106	0.440	1

(first two columns and two rows of the table).

We can observe a positive and relatively high correlation among the number of products, features and test suite size. This is somewhat expected because the number of valid products is expected to increase when more features are added to a feature model. In the same sense, more features not only usually imply more combinations of features that must be covered by the test suite, but also usually mean that more test cases must be added.

Regarding performance, we expect the algorithms to take more time to generate the test suites for larger models. The positive correlation between the performance and the three previous size-related measures (products, features and size) supports this idea. However, the value is too low (around 0.3 on average) to clearly claim that larger models require more computation time.

The correlation coefficient between the similarity metric and the other metrics is low except for the performance where the higher the similarity, the longer time spent in building the test suite. In the case of number of products, the value is rather small (-0.015) to draw any conclusions. In the case of test suite size, we would expect larger test suites to have higher similarity values (positive correlation with test suite size), but this is not the case in general. We believe these results might be due to the fact that the similarity metric, as defined in Equation (3), only considers selected variant features; however, we would argue that unselected features must also be considered in computing similarity because they are also part of the t-sets, which should be covered by the test suites. It is part of our future work to evaluate alternatives to this metric.

Let us now analyse the metrics results grouped by algorithms, shown in Table 4, for the 19 feature models. In the table we highlight with dark gray the results that are the best for each metric. The results of the benchmark models reveal that CASA is the best algorithm regarding the size of the test suite (with a statistically significant difference), whereas for PGS and ICPL the differences in test suite size are not statistically significant. If we focus on performance time, ICPL is the clear winner followed by CASA. PGS is outperformed by CASA in test suite size and computation time. PGS is also the slowest algorithm, in part because it is not specifically designed to deal with feature models. Regarding the similarity metric, ICPL is the algorithm providing more dissimilar products (with a statistically significant difference) and CASA is the second one, but with no statistically significant difference with PGS. This is, as we have mentioned before, somewhat counter-intuitive because ICPL produces larger test suites than CASA, therefore we would expect that test suites with more products to have

more similar products considering that the number of features in the feature models is finite. To elucidate why this is the case is part of our future work.

Table 4: Average of the metrics computed on the test suites generated by the three algorithms for the 19 feature models.

Model	Algor.	Size	Performance	Similarity
Apache	CASA	6.00	566.67	0.3635
	PGS	8.13	27196.40	0.3450
	ICPL	8.00	189.67	0.3044
argo-uml-spl	CASA	6.00	600.00	0.3670
	PGS	7.97	21167.77	0.3611
	ICPL	8.00	321.53	0.3028
BerkeleyDB	CASA	6.00	12600.00	0.3721
	PGS	8.00	126936.77	0.4478
	ICPL	7.00	2027.57	0.2679
BerkeleyDBF	CASA	6.00	533.33	0.3645
	PGS	8.10	28591.70	0.3515
	ICPL	7.00	176.77	0.3232
BerkeleyDBM	CASA	30.00	7333.33	0.3482
	PGS	30.70	249834.47	0.3597
	ICPL	31.00	554.47	0.2684
BerkeleyDBP	CASA	9.53	4133.33	0.3826
	PGS	11.43	65988.83	0.3915
	ICPL	10.00	366.03	0.3576
Curl	CASA	8.00	916.67	0.3537
	PGS	12.13	43605.03	0.3490
	ICPL	12.00	276.03	0.2634
DesktopSearcher	CASA	8.37	2266.67	0.3785
	PGS	9.20	32552.70	0.3951
	ICPL	9.07	412.20	0.3895
fame-dbms-fm	CASA	10.00	1700.00	0.3255
	PGS	13.80	58227.27	0.3327
	ICPL	11.00	378.00	0.3190
gpl	CASA	12.00	1966.67	0.3494
	PGS	13.13	62859.50	0.3576
	ICPL	13.00	317.83	0.3481
LinkedList	CASA	12.13	2133.33	0.4057
	PGS	15.43	74601.10	0.4151
	ICPL	14.00	462.53	0.3988
LLVM	CASA	6.00	633.33	0.3653
	PGS	8.73	32615.13	0.3523
	ICPL	9.00	221.73	0.2320
PKJab	CASA	6.00	550.00	0.3752
	PGS	7.63	28318.90	0.3726
	ICPL	7.00	193.13	0.3439
Prevayler	CASA	6.00	550.00	0.3610
	PGS	6.60	26052.00	0.3598
	ICPL	8.00	156.20	0.2677
SensorNetwork	CASA	10.23	1583.33	0.3719
	PGS	12.30	63212.33	0.3595
	ICPL	14.00	445.27	0.3166
TankWar	CASA	12.50	39200.00	0.3483
	PGS	14.77	152199.10	0.3571
	ICPL	14.00	612.23	0.3140
Wget	CASA	9.00	766.67	0.3548
	PGS	12.43	46869.77	0.3541
	ICPL	12.00	290.87	0.2685
x264	CASA	16.00	2966.67	0.3523
	PGS	16.97	74500.10	0.3640
	ICPL	17.00	359.53	0.2574
ZipMe	CASA	6.00	533.33	0.3505
	PGS	7.43	26376.53	0.3748
	ICPL	7.00	165.40	0.3429

Finally, in Table 5 we show the  $\hat{A}_{12}$  statistic to assess the practical significance of the results. Given a performance measure  $M$ ,  $\hat{A}_{12}$  measures the probability that running algorithm  $A$  yields higher  $M$  values than running another algorithm  $B$ . If the results of the two algorithms are equal, then  $\hat{A}_{12} = 0.5$ . If  $\hat{A}_{12} = 0.3$  entails one would obtain higher values for  $M$  with algorithm  $A$ , 30% of the times. In Table 5 we have highlighted the largest distance from 0.5 (equality) per quality indicator, note that 0.5 indicates no difference in the comparison. Recall that we have to highlight two

	Size			Performance			Similarity		
	CASA	ICPL	PGS	CASA	ICPL	PGS	CASA	ICPL	PGS
CASA	-	0.3312	0.3194	-	0.9286	0.0109	-	0.8479	0.4807
ICPL	0.6688	-	0.4653	0.0714	-	0.0005	0.1521	-	0.1577
PGS	0.6806	0.5347	-	0.9891	0.9995	-	0.5193	0.8423	-

Table 5:  $\hat{A}_{12}$  statistical test results.

values per metric, because the direction of the comparison does not affect the interpretation of the result, although the value itself is complementary (both adding up to 1).

Regarding size, there is not statistically significant difference between ICPL and PGS, while CASA obtains the best results in more than 66% of the times. Regarding performance, ICPL is faster with a higher probability than the other algorithms. ICPL is faster than CASA in 93% of the times, moreover, ICPL is faster than PGS in 99.95% of the times. So, ICPL is clearly the best algorithm in performance without any doubts. Regarding similarity, ICPL is again the algorithm which obtains more dissimilar test suites. It obtains a lower value of similarity than CASA and PGS, in around 85% of the comparisons. As we have commented earlier in this section, this results of similarity are somehow unexpected, because smaller test suites ought to be more dissimilar than larger ones. For this reason, CASA would obtain lower values of test suite similarity, but it does not. Again, investigating why this is the case is part of our future work.

## 8. RELATED WORK

There exists substantial literature on SPL testing [9, 12, 14, 22]. However, to the best of our knowledge there are neither benchmarks nor frameworks for comparing approaches. In the area of Search-Based Software Engineering a major research focus has been software testing [11, 18], where there exists a plethora of articles that compare testing algorithms using different metrics. For example, Mansour *et al.* [26] compare five algorithms for regression testing using eight different metrics (including quantitative and qualitative criteria). Similarly, Uyar *et al.* [36] compare different metrics implemented as fitness functions to solve the problem of test input generation. To the best of our knowledge, in the literature on test case generation there is no well-known comparison framework for the research and practitioner community to use. Researchers usually apply their methods to open source programs and compute some metrics directly such as the success rate, the number of test cases and performance. The closest to a common comparison framework we could trace is the work of Rothermel and Harrold [32] where they propose a framework for regression testing.

## 9. CONCLUSIONS AND FUTURE WORK

In this research-in-progress paper, we put forward 19 feature models as a basis for a benchmark of CIT testing of SPLs. With this benchmark, we made an assessment of the comparison framework proposed by Perrouin *et al.* using three approaches (CASA, PGS and ICPL) for the case of pairwise testing. Overall the framework helped us identify facts such as that CASA obtains the smallest test suites, while ICPL is the fastest algorithm and also obtains the most dissimilar products. However, we also identified two shortcomings of this framework: *i)* similarity does not consider features that are not selected in a product, a fact that might skew the

expected output, and *ii)* tuple frequency is applicable on a per tuple basis only, so its value as an aggregative measure is not clear. As future work we plan to evaluate other metrics that could be used to complement the framework, for this we will follow the guidelines for metrics selection suggested in [27]. In addition, we expect to integrate more feature models into the benchmark as well as to refine or extend the feature model selection criteria.

## 10. ACKNOWLEDGEMENTS

This research is partially funded by the Austrian Science Fund (FWF) project P21321-N15 and Lise Meitner Fellowship M1421-N15, the Spanish Ministry of Economy and Competitiveness and FEDER under contract TIN2011-28194. We thank Martin Johansen and Øystein Haugen for their help with SPLCA tool, Norbert Siegmund for his support with SPLConqueror, and B.J. Garvin for his help with CASA.

## 11. APPENDIX

Let  $occurrence_p(ts, fs)$  compute the number of times (i.e. 0 or 1) that  $t$ -set  $ts$  appears in feature set  $fs$ . Thus we define  $occurrence(ts, tCA)$  as:

$$occurrence(ts, tCA) = \sum_{fs \in tCA} occurrence_p(ts, fs)$$

**Theorem 1.** The average tuple frequency depends only on the number of features,  $|FL|$ , and valid  $t$ -sets,  $|TS|$ .

*Proof.* We can write the average tuple frequency as:

$$\begin{aligned} & \frac{1}{|TS|} \sum_{ts \in TS} \frac{occurrence(ts, tCA)}{|tCA|} \\ &= \frac{1}{|TS| \cdot |tCA|} \sum_{ts \in TS} occurrence(ts, tCA) \\ &= \frac{1}{|TS| \cdot |tCA|} \sum_{ts \in TS} \sum_{fs \in tCA} occurrence_p(ts, fs) \\ &= \frac{1}{|TS| \cdot |tCA|} \sum_{fs \in tCA} \left( \sum_{ts \in TS} occurrence_p(ts, fs) \right) \\ &= \frac{1}{|TS| \cdot |tCA|} \sum_{fs \in tCA} \frac{|FL|(|FL| - 1)}{2} \\ &= \frac{|tCA|}{|TS| \cdot |tCA|} \frac{|FL|(|FL| - 1)}{2} \\ &= \frac{|FL|(|FL| - 1)}{2|TS|} \end{aligned}$$

Note the expression within the parentheses. This is the number of valid tuples in a feature set  $fs$ . Let us now select two arbitrary features from  $FL$ . These features can be both selected in  $fs$ , both unselected or one selected and the other not. In any case, as the product is a valid product, there exists a valid  $t$ -set in  $TS$  having these two arbitrary features that is covered by  $fs$ . Thus, the sum within the parentheses is the number of pairs of features:

$$\sum_{ts \in TS} occurrence_p(ts, fs) = \frac{|FL|(|FL| - 1)}{2}$$

□

## 12. REFERENCES

- [1] Argouml-spl project, 2013. <http://argouml-spl.tigris.org/>.
- [2] FeatureHouse website, 2013. <http://www.fosd.de/fh>.
- [3] Software Product Line Online Tools(SPLOOT), 2013. <http://www.splot-research.org/>.
- [4] SPL2go website, 2013. <http://spl2go.cs.ovgu.de/>.
- [5] A. Arcuri and L. Briand. A hitchhiker’s guide to statistical tests for assessing randomized algorithms in software engineering. *Softw. Test. Verif. Reliab*, 2012. in press.
- [6] D. Benavides, S. Segura, and A. R. Cortés. Automated analysis of feature models 20 years later: A literature review. *Inf. Syst.*, 35(6):615–636, 2010.
- [7] H. Cichos, S. Oster, M. Lochau, and A. Schürr. Model-based coverage-driven test suite generation for software product lines. In J. Whittle, T. Clark, and T. Kühne, editors, *MoDELS*, volume 6981 of *Lecture Notes in Computer Science*, pages 425–439. Springer, 2011.
- [8] M. B. Cohen, M. B. Dwyer, and J. Shi. Constructing interaction test suites for highly-configurable systems in the presence of constraints: A greedy approach. *IEEE Trans. Software Eng.*, 34(5):633–650, 2008.
- [9] P. A. da Mota Silveira Neto, I. do Carmo Machado, J. D. McGregor, E. S. de Almeida, and S. R. de Lemos Meira. A systematic mapping study of software product lines testing. *Information & Software Technology*, 53(5):407–423, 2011.
- [10] E. S. de Almeida, C. Schwanninger, and D. Benavides, editors. *16th International Software Product Line Conference, SPLC ’12, Salvador, Brazil - September 2-7, 2012, Volume 1*. ACM, 2012.
- [11] F. G. de Freitas and J. T. de Souza. Ten years of search based software engineering: A bibliometric analysis. In M. B. Cohen and M. Ó. Cinnéide, editors, *SSBSE*, volume 6956 of *Lecture Notes in Computer Science*, pages 18–32. Springer, 2011.
- [12] I. do Carmo Machado, J. D. McGregor, and E. S. de Almeida. Strategies for testing products in software product lines. *ACM SIGSOFT Software Engineering Notes*, 37(6):1–8, 2012.
- [13] J. J. Durillo and A. J. Nebro. jMetal: A java framework for multi-objective optimization. *Advances in Engineering Software*, 42(10):760 – 771, 2011.
- [14] E. Engström and P. Runeson. Software product line testing - a systematic mapping study. *Information & Software Technology*, 53(1):2–13, 2011.
- [15] F. Ensan, E. Bagheri, and D. Gasevic. Evolutionary search-based test generation for software product line feature models. In J. Ralyté, X. Franch, S. Brinkkemper, and S. Wrycza, editors, *CAiSE*, volume 7328 of *Lecture Notes in Computer Science*, pages 613–628. Springer, 2012.
- [16] J. Ferrer, P. M. Kruse, J. F. Chicano, and E. Alba. Evolutionary algorithm for prioritized pairwise test data generation. In T. Soule and J. H. Moore, editors, *GECCO*, pages 1213–1220. ACM, 2012.
- [17] B. J. Garvin, M. B. Cohen, and M. B. Dwyer. Evaluating improvements to a meta-heuristic search for constrained interaction testing. *Empirical Software Engineering*, 16(1):61–102, 2011.
- [18] M. Harman, S. A. Mansouri, and Y. Zhang. Search-based software engineering: Trends, techniques and applications. *ACM Comput. Surv.*, 45(1):11, 2012.
- [19] A. Hervieu, B. Baudry, and A. Gotlieb. Pacogen: Automatic generation of pairwise test configurations from feature models. In T. Dohi and B. Cukic, editors, *ISSRE*, pages 120–129. IEEE, 2011.
- [20] M. F. Johansen, Ø. Haugen, and F. Fleurey. An algorithm for generating t-wise covering arrays from large feature models. In de Almeida et al. [10], pages 46–55.
- [21] K. Kang, S. Cohen, J. Hess, W. Novak, and A. Peterson. Feature-Oriented Domain Analysis (FODA) Feasibility Study. Technical Report CMU/SEI-90-TR-21, Software Engineering Institute, Carnegie Mellon University, 1990.
- [22] J. Lee, S. Kang, and D. Lee. A survey on software product line testing. In de Almeida et al. [10], pages 31–40.
- [23] M. Lochau, S. Oster, U. Goltz, and A. Schürr. Model-based pairwise testing for feature interaction coverage in software product line engineering. *Software Quality Journal*, 20(3-4):567–604, 2012.
- [24] R. E. Lopez-Herrejon and D. S. Batory. A standard problem for evaluating product-line methodologies. In J. Bosch, editor, *GCSE*, volume 2186 of *Lecture Notes in Computer Science*, pages 10–24. Springer, 2001.
- [25] R. E. Lopez-Herrejon, J. Ferrer, F. Chicano, E. N. Haslinger, A. Egyed, and E. Alba. Paper code and data repository, 2014. <http://neo.lcc.uma.es/staff/javi/resources.html>.
- [26] N. Mansour, R. Bahsoon, and G. Baradhi. Empirical comparison of regression test selection algorithms. *Journal of Systems and Software*, 57:79–90, 2001.
- [27] A. Meneely, B. H. Smith, and L. Williams. Validating software metrics: A spectrum of philosophies. *ACM Trans. Softw. Eng. Methodol.*, 21(4):24, 2012.
- [28] S. Oster, F. Markert, and P. Ritter. Automated incremental pairwise testing of software product lines. In J. Bosch and J. Lee, editors, *SPLC*, volume 6287 of *Lecture Notes in Computer Science*, pages 196–210. Springer, 2010.
- [29] G. Perrouin, S. Oster, S. Sen, J. Klein, B. Baudry, and Y. L. Traon. Pairwise testing for software product lines: comparison of two approaches. *Software Quality Journal*, 20(3-4):605–643, 2012.
- [30] G. Perrouin, S. Sen, J. Klein, B. Baudry, and Y. L. Traon. Automated and scalable t-wise test case generation strategies for software product lines. In *ICST*, pages 459–468. IEEE Computer Society, 2010.
- [31] K. Pohl, G. Bockle, and F. J. van der Linden. *Software Product Line Engineering: Foundations, Principles and Techniques*. Springer, 2005.
- [32] G. Rothermel and M. J. Harrold. A framework for evaluating regression test selection techniques. In *Proceedings of the 16th International Conference on Software Engineering*, pages 201–210, 1994.
- [33] D. J. Sheskin. *Handbook of Parametric and Nonparametric Statistical Procedures*. Chapman & Hall/CRC; 4 edition, 2007.



- [34] N. Siegmund, M. Rosenmüller, C. Kästner, P. G. Giarrusso, S. Apel, and S. S. Kolesnikov. Scalable prediction of non-functional properties in software product lines: Footprint and memory consumption. *Information & Software Technology*, 55(3):491–507, 2013.
- [35] M. Svahnberg, J. van Gurp, and J. Bosch. A taxonomy of variability realization techniques. *Softw., Pract. Exper.*, 35(8):705–754, 2005.
- [36] H. T. Uyar, A. S. Uyar, and E. Harmanci. Pairwise sequence comparison for fitness evaluation in evolutionary structural software testing. In *Genetic and Evolutionary Computation Conference*, pages 1959–1960, 2006.
- [37] A. Vargha and H. D. Delaney. A critique and improvement of the cl common language effect size statistics of mcgraw and wong. *Journal of Educational and Behavioral Statistics*, 25(2):101–132, 2000.